Visualising rate of change: application to age-specific fertility

Han Lin Shang[1]

Australian National University

arXiv:1406.4386v2 [stat.AP] 9 Dec 2015

Abstract

Visualisation methods help in the discovery of characteristics that might not have been apparent using mathematical models and summary statistics, and yet it has not received much attention in demography with the exceptions of scatter plot and Lexis surface. We introduce phase-plane plot as a means of visualising the rate of change, obtained from the first and second derivatives of a continuous function. Illustrated by the age-specific Australian fertility rates, we present two phase-plane plots for years 1921 and 2006 and highlight their differences. Similar to the scatter plot and Lexis surface, the phase-plane plot identifies the age associated with maximum fertility and displays skewness of age-specific fertility distribution. Differing from the scatter plot and Lexis surface, the phase-plane plot identifies the age associated with maximum positive or negative velocity, is able to compare the magnitude of the rate of change between any two years based on size of radius.

Keywords: B-spline basis functions; Data visualisation; Functional data; Velocity; Acceleration

[1]Running title: Phase-plane plot

Postal address: Research School of Finance, Actuarial Studies and Statistics, Level 4, Building 26C, Kingsley Street, Australian National University, Acton, ACT 2601, Australia

Email: hanlin.shang@anu.edu.au

Introduction

Methods for analysing age-specific demographic data can be loosely divided into two categories: those considering age either as a discrete variable in real space R^p (p represents the total number of ages) or as a continuous variable in function space bounded within a finite interval. In the discrete case, Bell (1992), Lee (1993), Lee and Carter (1992), and Li and Lee (2005) have done extensive work on modelling and forecasting age-specific fertility and mortality, where the demographic rates at different ages are considered as discrete data points. In the continuous case, D'Amato et al. (2011), D'Amato et al. (2012), Hyndman et al. (2013), and Hyndman and Ullah (2007) introduced functional data analysis to model and forecast age-specific demographic rates at a given year as a continuous and smooth function, which can later be converted to discrete ages at any sampling interval. Hyndman and Ullah (2007) emphasised that the main advantage of functional data analysis is that the smoothing technique can be incorporated into the modelling and forecasting of functions, so any missing or noisy data can be naturally dealt with.

Aside from incorporating smoothness into modelling and forecasting, the ability to consider derivatives, a by-product of conceiving of the data as functions, is also of great advantage for visualisation. The first derivative, also known as velocity in physics, measures the rate at which the value of a function $f(x)$ changes with respect to the change of a variable x. The second derivative, also known as acceleration, measures the rate at which the rate of change alters with respect to the change of variable. For example, Ramsay and Silverman (2002 Chapter 6) considered human growth curve. When a child is growing more quickly, the velocity increases rapidly, decreasing when growth slows. When a child starts to grow, the absolute acceleration is at its maximum, before gradually decreasing during the rapid growing period. Then, the absolute acceleration starts to increase at the slow growing period, before stabilises when a child stops growing. Thus, analysing velocity and acceleration information may reveal patterns that are not easily seen from the original functions themselves. Illustrated by Australian age-specific fertility rates in years 1921 and 2006, we introduce the phase-plane plot for visualising the rate of change. More information on some properties of B-spline basis functions are deferred

to Appendix, and computational code in R are available in the supplementary material.

Data set

We consider annual age-specific Australian fertility rates from 1921 to 2006. The data set has been obtained from the Australian Bureau of Statistics (Cat. No. 3105.0.65.001, Table 38) and is available in the *rainbow* package (Shang and Hyndman 2015) in R (R Core Team 2015). The data consist of annual fertility rates by age of mother in single years from ages 15 to 49. Graphical displays of the data are given in Figure 1. From the rainbow plot in Figure 1(a), we see the phenomenon of fertility postponement in the most recent years shown in purple. From the Lexis surface in Figure 1(b), we see the increases in fertility between ages 20 and 30 from 1940 to 1980, this reflects the baby boom period occurred around 1960.

[Figure 1 about here.]

Visualising rate of change: phase-plane plot

Derivatives are useful in visualising functional data. To visualise the rate of change in age-specific fertility, we consider the phase-plane plot in functional data analysis (Ramsay and Silverman 2002 Section 3.3). It is a two-dimensional plot of the acceleration against the velocity of a function. Let us consider a simple function $\sin(2\pi x)$, which has the first derivative $2\pi\cos(2\pi x)$ and second derivative $-4\pi\sin(2\pi x)$, for $x = 0, 0.01, \ldots, 1$. A two-dimensional plot of the acceleration against the velocity is shown in Figure 2, where the highest absolute velocity or acceleration is reached when the other variable is zero.

[Figure 2 about here.]

From the highest absolute acceleration to the highest absolute velocity (1st and 3rd quadrants), it displays an increase in rate of change. From the highest absolute velocity to the highest absolute acceleration (2nd and 4th quadrants), it shows the rate of change goes to zero. When the velocity reaches zero, it indicates the occurrence of maximum or

minimum fertility. As noted by Ramsay and Silverman (2006 pp.30-33), the size of the radius can imply the magnitude of rate of change; the greater the radius is, the greater the magnitude of rate of change.

The phase-plane plot allows us to visualise changes in the shapes of the cycles from year to year, from which we aim to reveal social changes in age-specific fertility, including the phenomena of baby boom occurred around 1960 and postponement in recent years. In Figure 3, we display the phase-plane plot and scatter plot for age-specific fertility at the first year of data in 1921 and last year of data in 2006 to highlight their contrasts in age distribution of fertility rates.

[Figure 3 about here.]

From Figure 3(a), the age-specific fertility rates in 1921 have an increase in velocity from ages 15 to 19, they reach the highest positive velocity at 19 before a decreasing velocity from 20 to 26. The maximum fertility occurs at 26, as shown by zero velocity. The absolute velocity increases slowly from 27 to 40 and it reaches the maximum negative velocity at 41. The absolute velocity then decreases gradually to zero from 42 to 49. Since the horizontal location of the centre is to the right, the age distribution of fertility rates is right-skewed. It indicates a net positive velocity, where the speed increasing to the maximum fertility is faster than the speed decreasing to the minimum fertility.

By contrast, the age-specific fertility rates in 2006 show an increase in velocity from 15 to 18, then a decrease from 19 to 21 before another increase in velocity from 22 to 26. This indicates the presence of bimodality in data. The maximum positive velocity is reached at 26, then the velocity decreases from 27 to 30. The maximum fertility occurs at about 31, as shown by zero velocity. From 32 to 35, there is an increase in absolute velocity, before reaching the maximum negative velocity at 36. From 37 to 49, the absolute velocity gradually goes to zero. Since the horizontal location of the centre is to the left, the age distribution of fertility rates is left-skewed. It indicates a net negative velocity, where the speed increasing to the maximum fertility is slower than the speed decreasing to the minimum fertility.

Conclusion

From an aspect of functional data analysis, we consider the problem of visualisation in demography. Acknowledging the fact that the phase-plane plot can be applied to visualise rate of change in the Australian age-specific fertility from 1921 to 2006, we introduce the phase-plane plot to visualise the rate of change in years 1921 and 2006. Similar to the scatter plot and Lexis surface, the phase-plane plot identifies the age associated with the maximum fertility and it displays the skewness of age-specific fertility distribution based on the net velocity. Differing from the scatter plot and Lexis surface, the phase-plane plot identifies the age associated with maximum positive or maximum negative velocity and compare the magnitude of the rate of change between any two years based on the size of the radius. Thus, the phase-plane plot should be considered as a complementary graphical tool to the scatter plot and Lexis surface for exploring rate of change in age-specific demographic data, including age-specific mortality, migration and population size.

Acknowledgments

The author extends sincere thanks to Peter W. F. Smith for insightful discussion during an early stage of this paper.

Appendix: B-spline basis functions and their derivatives

de Boor (2001) and Dierckx (1995) presented the authoritative account on B-spline basis functions. Here, we present the procedure for computing derivatives using B-spline basis functions. Let u be a set of $m+1$ non-decreasing sequences $u_0 \leq u_1 \leq u_2 \leq \cdots \leq u_m$. These u_is are known as knots and the finite interval $[u_i, u_{i+1})$ is known as the i^{th} knot span for $i = 0, \ldots, m-1$. Since B-spline basis functions are a type of piecewise polynomials, the method requires the specification of degree p. The i^{th} B-spline basis function of degree p, written as $N_{i,p}(u)$ can be defined recursively as follows:

$$N_{i,0}(u) = \begin{cases} 1 & \text{if } u_i \leq u < u_{i+1}; \\ 0 & \text{if otherwise.} \end{cases}$$

$$N_{i,p}(u) = w_{i,p-1} \times N_{i,p-1}(u) + (1 - w_{i+1,p-1}) \times N_{i+1,p-1}(u), \quad \text{for } p \geq 1, \tag{1}$$

where $w_{i,p-1} = (u - u_i)/(u_{i+p} - u_i)$. The above recursion is referred to as the Cox–de Boor recursion formula (see Cox 1972; de Boor 1972 p.90). When $p = 0$, it corresponds to the step function, assigning 1 if u is in the i^{th} knot span $[u_i, u_{i+1})$. When $p = 1$, $N_{i,1}(u)$ spans two intervals and is piecewise linear, $N_{i,2}(u)$ spans three intervals that is piecewise quadratic, $N_{i,3}(u)$ spans four intervals that is piecewise cubic. The cubic B-spline is the common choice, as it provides sufficient flexibility for approximating most functions.

A function can accurately be approximated by a sufficient set of B-spline basis functions and their corresponding coefficients, given by:

$$\sum_{i=0}^{m-1} c_i N_{i,p}(u), \tag{2}$$

where c_i represents the B-spline coefficients. An advantage of the B-spline is the easy computation of derivatives,

$$D\left(\sum_{i=0}^{m-1} c_i N_{i,p}(u)\right) = (p-1) \sum_{i=0}^{m} \frac{c_i - c_{i-1}}{u_{i+p-1} - u_i} N_{i,p-1}(u), \quad \text{for } p \geq 1, \tag{3}$$

where D represents the first derivative with respect to u and $c_{-1} := 0 =: c_m$. As noted by de Boor (2001), the first derivative of a spline function can be represented by differencing the B-spline coefficients and a spline of one order lower. The higher-order derivatives can be calculated iteratively.

As a demonstration, Figure 4(a) shows the seven cubic B-spline basis functions corresponding to three equally spaced interior knots over the interval $[-1, 1]$, whereas Figure 4(b) presents first derivative of the B-spline basis functions. From Figure 4(a), apart from the two end basis functions, each basis function begins at zero and at a certain knot position, rises to a peak before falling back to zero and remaining there until the right boundary. The first and last basis functions rise from the first and last interior knot to a value of one on the right and left boundary.

[Figure 4 about here.]

References

Bell, W. (1992), ARIMA and principal components models in forecasting age-specific fertility, *in* N. Keilman and H. Cruijsen, eds, 'National Population Forecasting in Industrialized Countries', Swets & Zeitlinger, Amsterdam, pp. 177–200.

Cox, M. G. (1972), 'The numerical evaluation of B-splines', *Journal of the Institute of Mathematics and its Applications* ***10***(2), 134–149.

D'Amato, V., Piscopo, G. and Russolillo, M. (2011), 'The mortality of the Italian population: smoothing technique on the Lee-Carter model', *The Annals of Applied Statistics* **5**(2A), 705–724.

D'Amato, V., Piscopo, G. and Russolillo, M. (2012), Simulation framework in fertility projections, *in* B. Apolloni, S. Bassis, A. Esposito and F. C. Morabito, eds, 'Neural Nets and Surroundings', Springer, Salerno, Italy, pp. 209–216.

de Boor, C. (1972), 'On calculating with B-splines', *Journal of Approximation Theory* **6**(1), 50–62.

de Boor, C. (2001), *A Practical Guide to Splines*, Vol. 27 of *Applied Mathematical Sciences*, revised edn, Springer, New York.

Dierckx, P. (1995), *Curve and Surface Fitting with Splines*, Monographs on Numerical Analysis, Oxford Science Publications, New York.

Hyndman, R. J., Booth, H. and Yasmeen, F. (2013), 'Coherent mortality forecasting: the product-ratio method with functional time series models', *Demography* **50**(1), 261–283.

Hyndman, R. J. and Ullah, M. S. (2007), 'Robust forecasting of mortality and fertility rates: a functional data approach', *Computational Statistics & Data Analysis* **51**(10), 4942–4956.

Lee, R. D. (1993), 'Modeling and forecasting the time series of US fertility: age distribution, range and ultimate level', *International Journal of Forecasting* **9**(2), 187–202.

Lee, R. D. and Carter, L. R. (1992), 'Modelling and forecasting U.S. mortality', *Journal of the American Statistical Association* **87**(419), 659–671.

Li, N. and Lee, R. (2005), 'Coherent mortality forecasts for a group of populations: an extension of the Lee-Carter method', *Demography* **42**(3), 575–594.

R Core Team (2015), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: http://www.R-project.org/

Ramsay, J. O. and Silverman, B. W. (2002), *Applied Functional Data Analysis*, Springer-Verlag, New York.

Ramsay, J. O. and Silverman, B. W. (2006), *Functional Data Analysis*, 2nd edn, Springer Science+Business Media LLC, New York.

Shang, H. L. and Hyndman, R. J. (2015), *rainbow: Rainbow plots, bagplots and boxplots for functional data*, R package version 3.3.
URL: http://cran.r-project.org/web/packages/rainbow

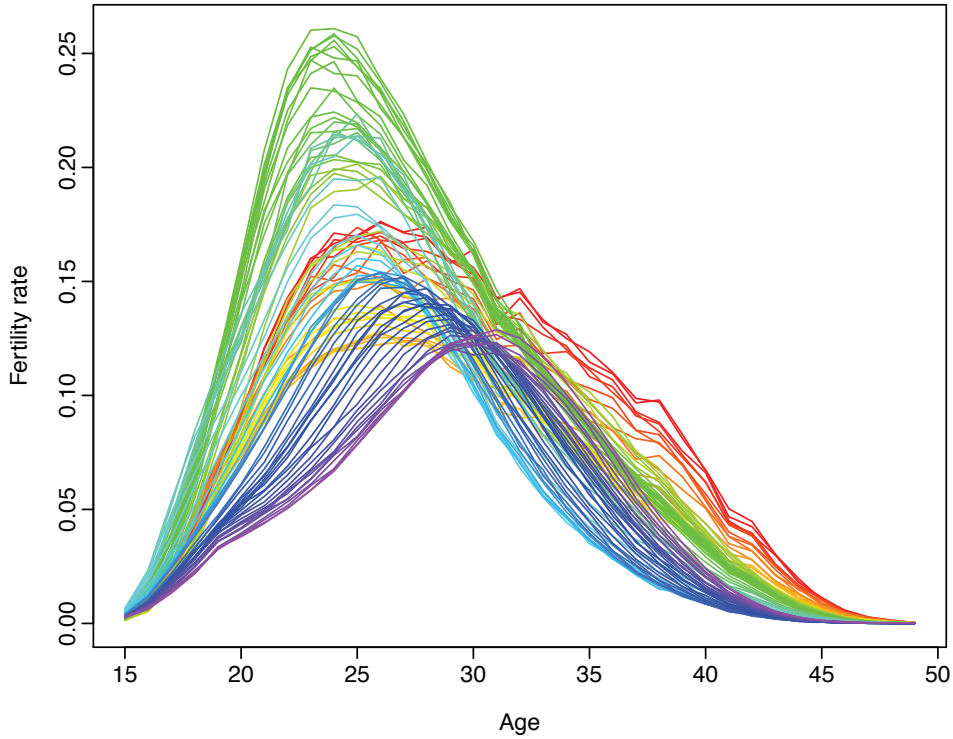

(a) Rainbow plot

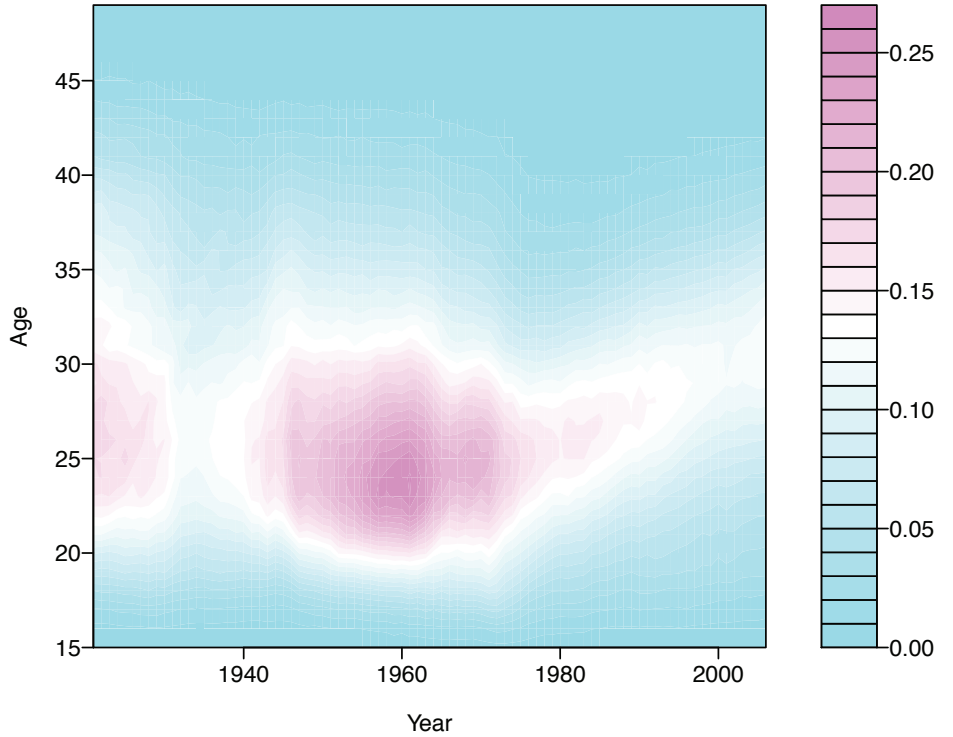

(b) Lexis surface

Figure 1: Plots of observed age-specific fertility rates for Australia from 1921 to 2006. Source: Australian Bureau of Statistics (Cat. No. 3105.0.65.001, Table 38).

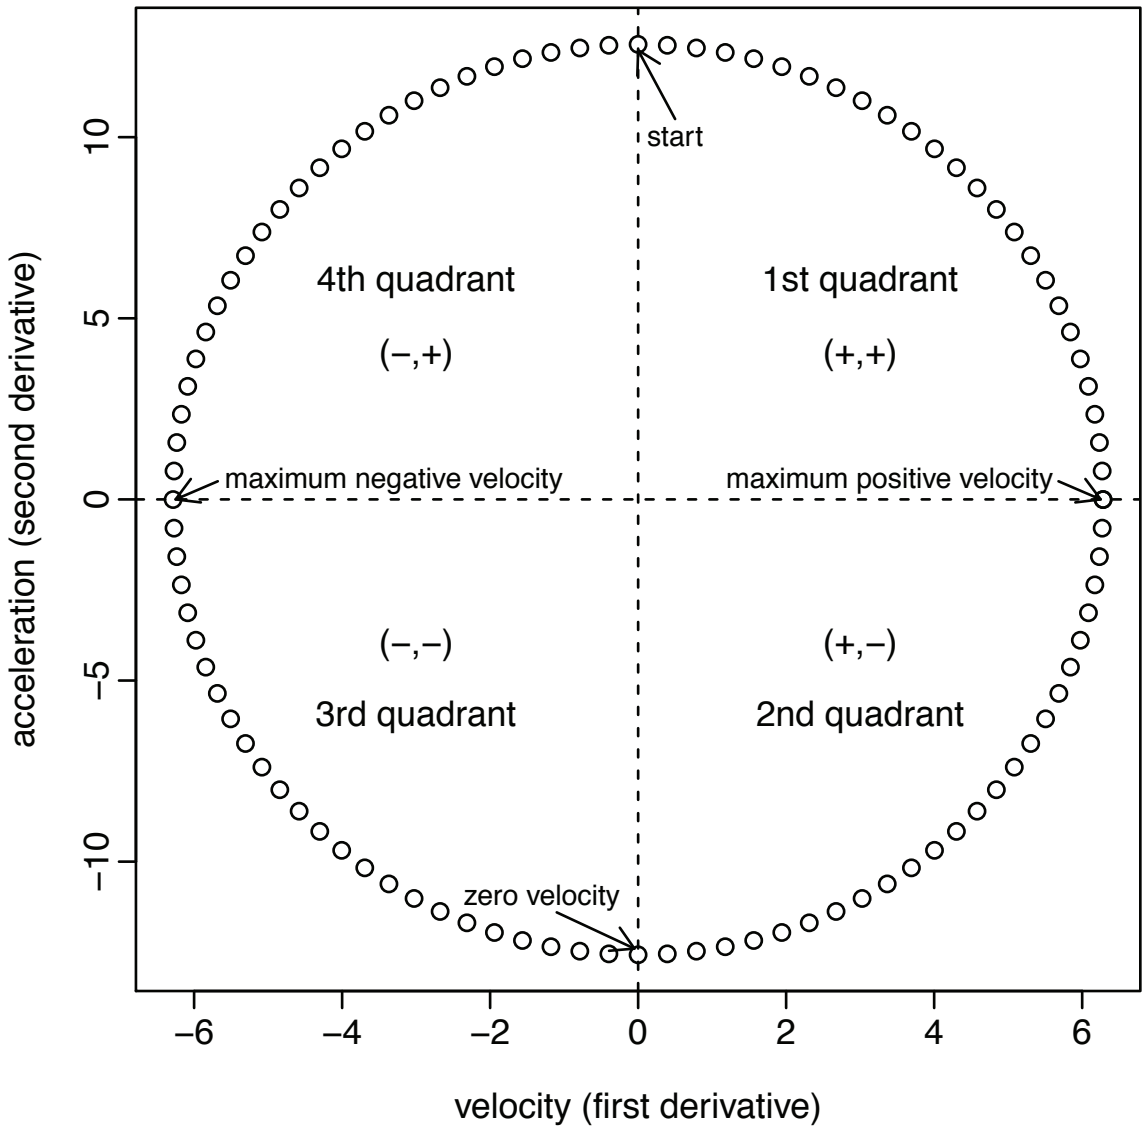

Figure 2: A two-dimensional plot of the acceleration against the velocity.

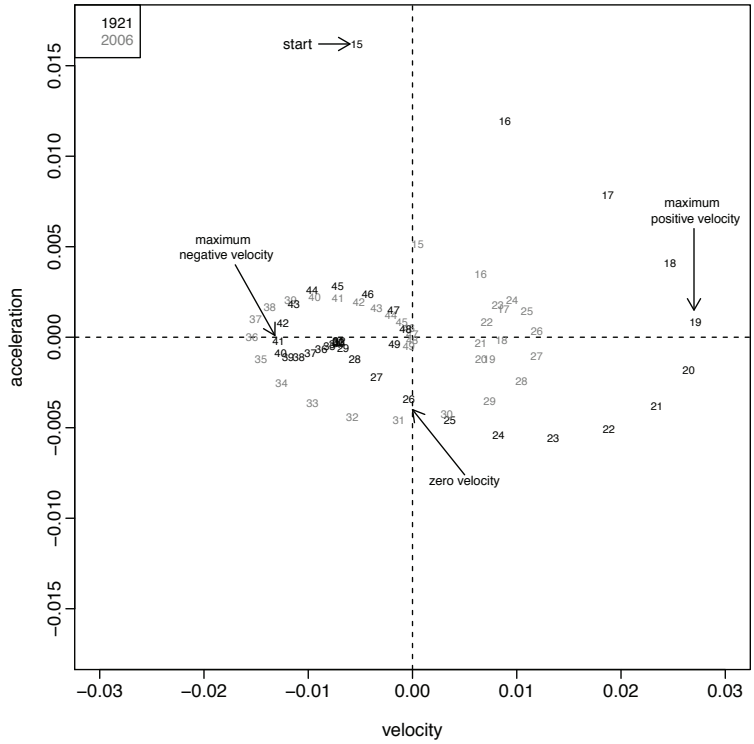

(a) Phase-plane plot

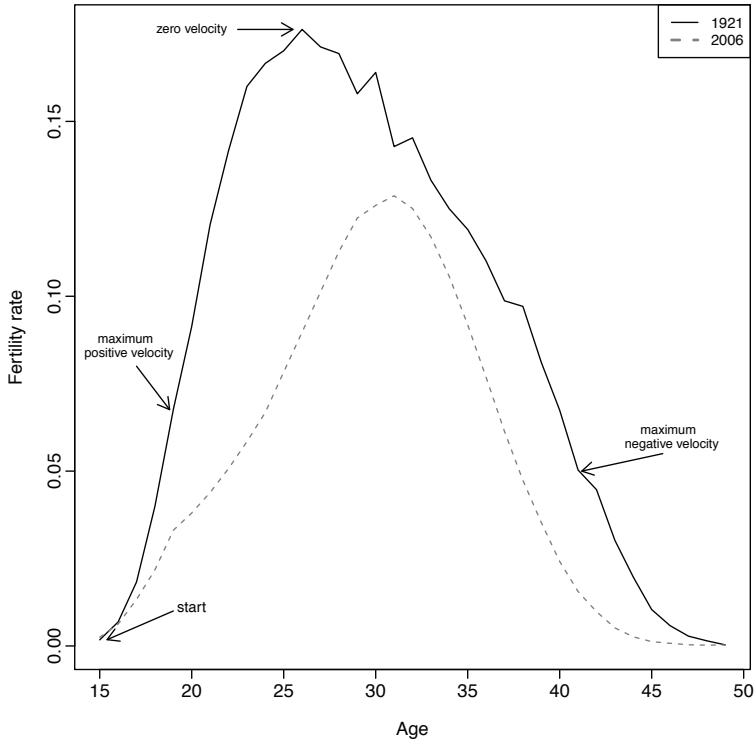

(b) Scatter plot

Figure 3: Phase-plane plot and scatter plot of the age-specific fertility rates for Australia in 1921 (as shown by dark grey colour) and 2006 (as shown by light grey colour).

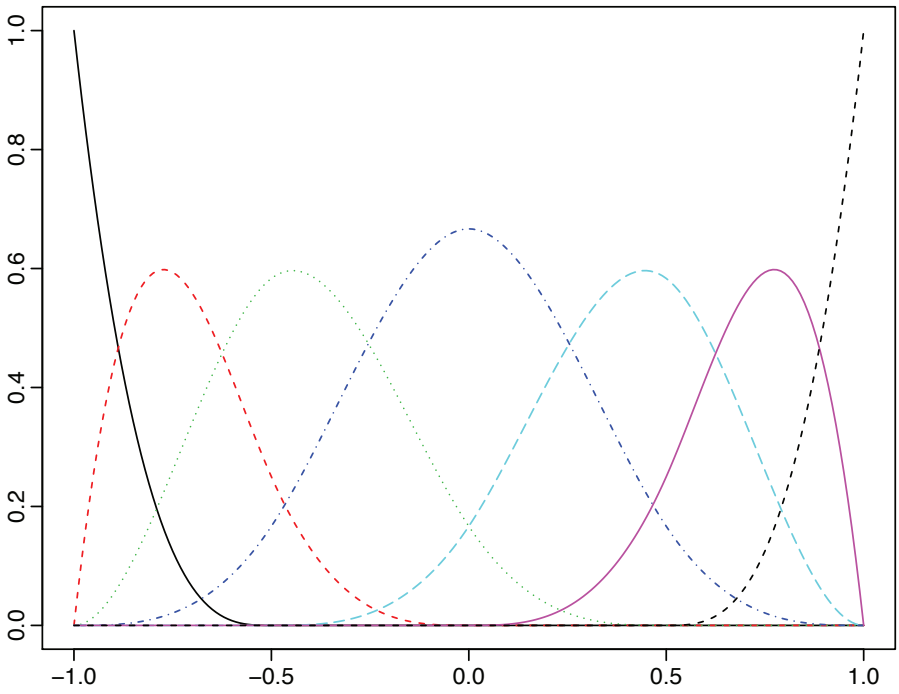

(a) B-spline basis functions

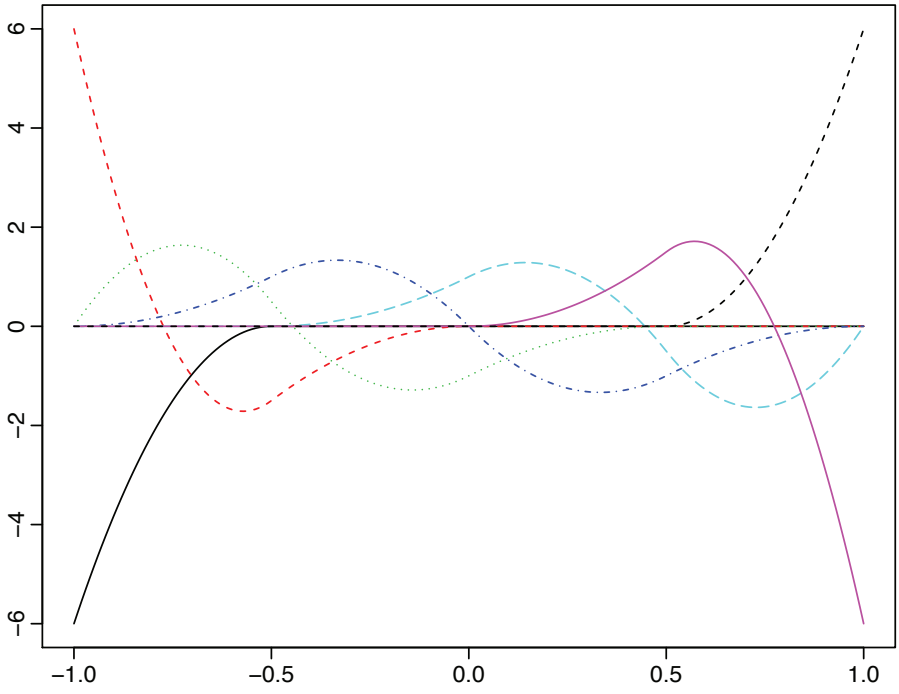

(b) First derivative of B-spline basis functions

Figure 4: For instance, we present the seven cubic B-spline basis functions defined over the interval $[-1, 1]$ by three interior knots.